

The Evolution of Cyber Infrastructure for Science and Engineering

William E. Johnston

<http://www-itg.lbl.gov/~wej/>

***Computational Research Division,
DOE Lawrence Berkeley National Lab***

and

***NASA Advanced Supercomputing (NAS) Division,
NASA Ames Research Center***

The Process of Large-Scale Science is Changing

- Data, computing, people, instruments are all distributed
- Instruments need to be integrated with large-scale computing and data systems
- The volume and complexity of data that must be dealt with is increasing dramatically
- In some fields – e.g. astronomy and astrophysics – science is being done by mining data from dozens of instruments, instead of the direct use of a single instrument like a telescope

The Process of Large-Scale Science is Changing

- Large-scale science and engineering problems *require collaborative use of many compute and data resources*, including supercomputers and large-scale data storage systems, *all of which must be integrated with applications and data* that are
 - developed by different teams of researchers
 - or that are obtained from different instrumentsand all of which are at different geographic locations. See*, e.g., refs.1.

(* See the last slide for references.)

Drivers for Science Cyberinfrastructure

- In order to inform the development and deployment of technology, *a set of high-impact science applications in the areas of high energy physics, climate, chemical sciences, and magnetic fusion energy have been analyzed** to characterize their visions for the future process of science, and the networking and middleware capabilities needed to support those visions

*DOE Office of Science, High Performance Network Planning Workshop. August 13-15, 2002: Reston, Virginia, USA.
<http://doecollaboratory.pnl.gov/meetings/hpnpw>

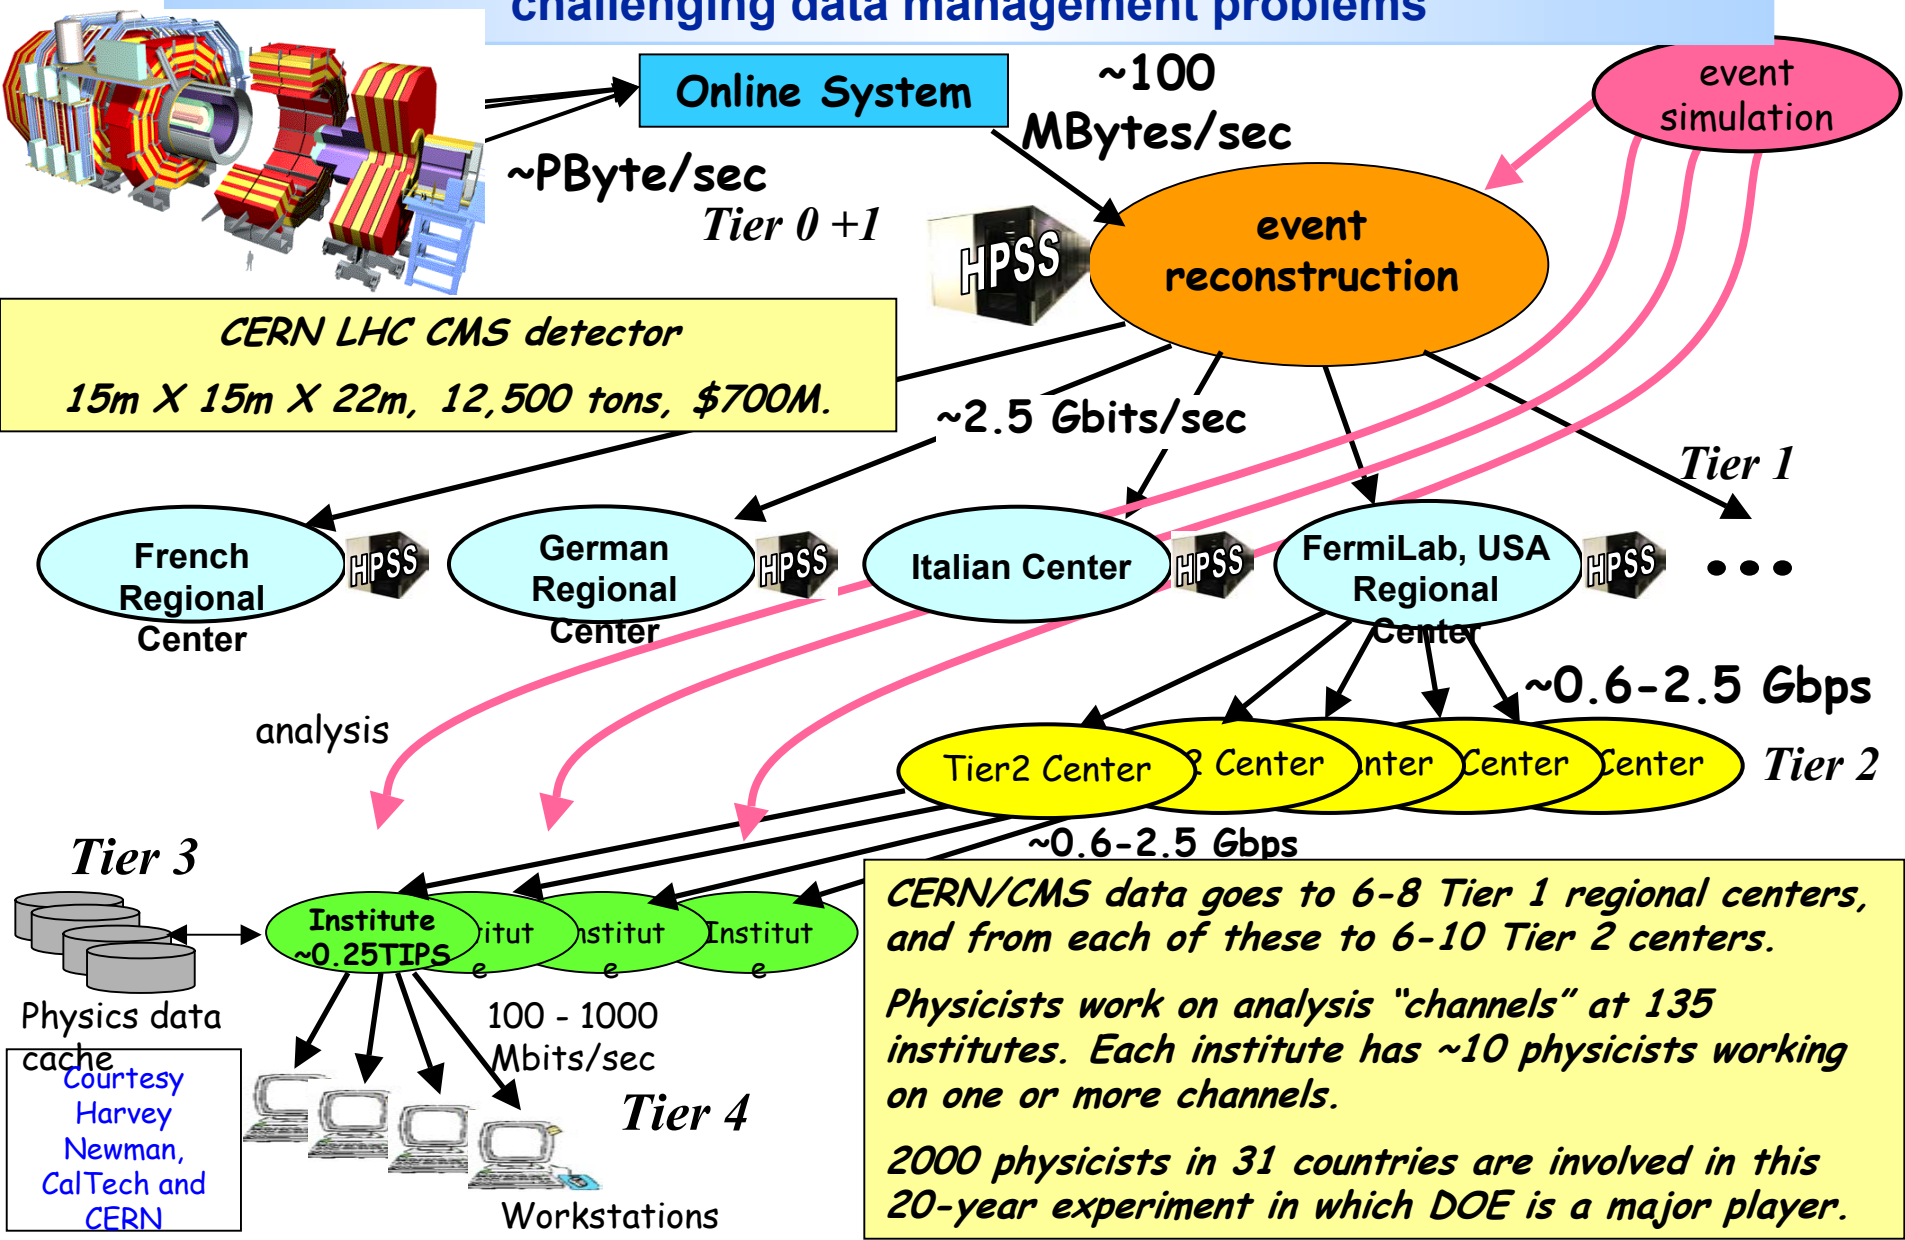
Visions for Future Science Process

	Feature	Characteristics that Motivate H-S Nets	Vision for the Future Process of Science	Anticipated Requirements	
Discipline				Networking	Middleware / Grid Services
Climate (near term)		<ul style="list-style-type: none"> • A few data repositories, many distributed computing sites • NCAR - 20 TBy • NERSC - 40 TBy • ORNL - 40 TBy 		<ul style="list-style-type: none"> • Authenticated data streams for easier site access through firewalls 	<ul style="list-style-type: none"> • Server side data processing (computing and cache embedded in the net) • Information servers for global data catalogues
Climate (5 yr)		<ul style="list-style-type: none"> • Add many simulation elements/components as understanding increases • 100 TBy / 100 yr generated simulation data, 1-5 PBy / yr (just at NCAR) <ul style="list-style-type: none"> – Distribute to major users in large chunks for post-simulation analysis 	<ul style="list-style-type: none"> • Enable the analysis of model data by all of the collaborating community 	<ul style="list-style-type: none"> • Robust access to large quantities of data 	<ul style="list-style-type: none"> • Reliable data/file transfer <ul style="list-style-type: none"> – Across system / network failures
Climate (5+ yr)		<ul style="list-style-type: none"> • 5-10 PBy/yr (at NCAR) • Add many diverse simulation elements/components, including from other disciplines - this must be done with distributed, multidisciplinary simulation • Virtualized data to reduce storage load 	<ul style="list-style-type: none"> • Integrated climate simulation that includes all high-impact factors 	<ul style="list-style-type: none"> • Robust networks supporting distributed simulation - adequate bandwidth and latency for remote analysis and visualization of massive datasets 	<ul style="list-style-type: none"> • Quality of service guarantees for distributed, simulations • Virtual data catalogues and work planners for reconstituting the data on demand • Knowledge management to facilitate science-level, Grid based problem solving systems (the “Semantic Grid”)



Example: High Energy Physics Data Management

CERN / LHC Data: One of Science's most challenging data management problems



High Energy Physics Data Management Issues

- Distributing petabytes of data per year to hundreds of sites around the world for analysis
 - Reliable, wide-area, high-volume data management
 - Global naming, replication, and caching of datasets
 - Easily accessible pools of computing resources
- Virtual data catalogues and on-demand data generation
 - Some types of analysis are pre-defined and catalogued prior to generation - and then the data products are generated on demand when the virtual data catalogue is accessed
 - Sometimes regenerating derived data is faster and easier than trying to store and/or retrieve that data from remote repositories
 - For similar reasons this is also of great interest to the EOS (Earth Observing Satellite) community

HEP Data is Similar to the EOS Data Scenario

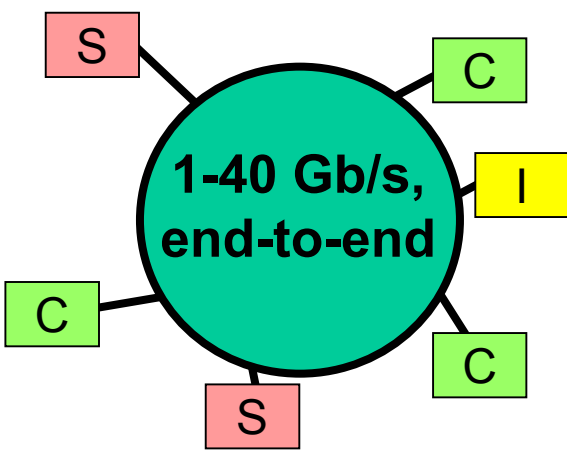
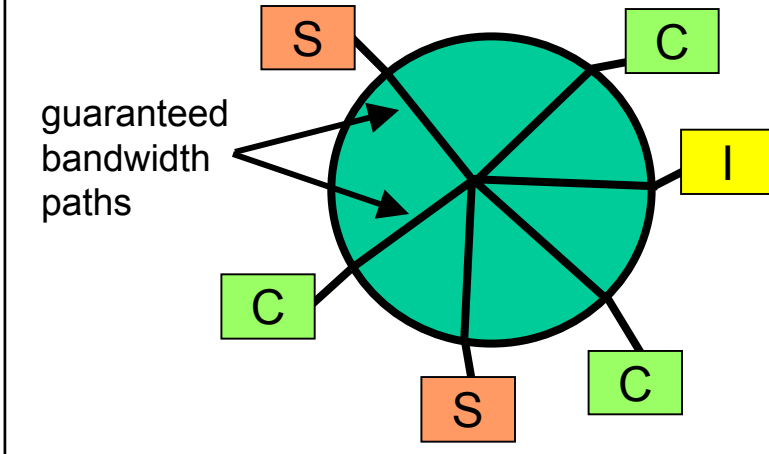
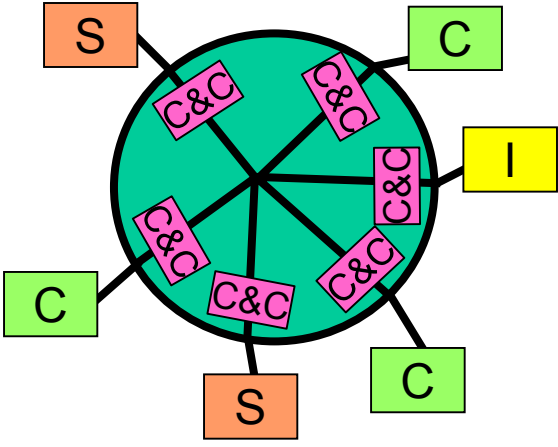
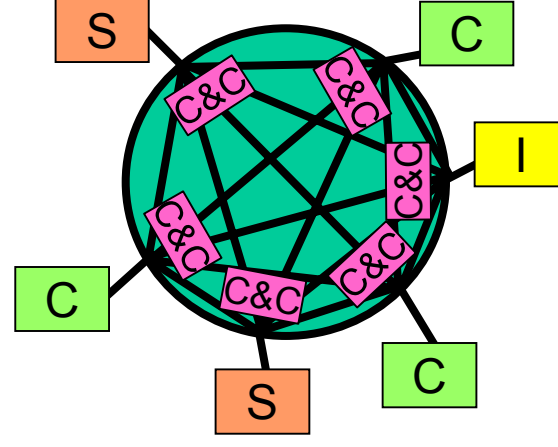
- The Earth Observing Satellite data scenario is the same scale of High Energy Physics, but with differences
 - EOS has multiple sources, HEP has a single source
 - Comparable data volumes
 - Comparable processing requirements
 - Different type of customer base (a few large customers – like physics) but also a lot of small customers (unlike physics)
 - EOS has a much greater diversity of data types

	Feature	Characteristics that Motivate H-S Nets	Vision for the Future Process of Science	Anticipated Requirements	
Discipline				Networking	Middleware / Grid Services
High Energy Physics (near term (1-2yr))		<ul style="list-style-type: none"> • Instrument based data sources • Hierarchical data repositories • Hundreds of analysis sites • Petabytes of data 	<ul style="list-style-type: none"> • Productivity aspects of rapid response 	<ul style="list-style-type: none"> • Gigabit/sec • end-to-end QoS 	<ul style="list-style-type: none"> • Secure access to world-wide resources • Data migration in response to usage patterns and network performance <ul style="list-style-type: none"> – naming and location transparency • Deadline scheduling for bulk transfers • Policy based scheduling / brokering for the ensemble of resources needed for a task • Automated planning and prediction to minimized time to complete task
HEP (3-5 yr)		<ul style="list-style-type: none"> • 100s of petabytes of data • Global collaboration • Compute and storage requirements will be satisfied by optimal use of all available resources 	<ul style="list-style-type: none"> • Worldwide collaboration will cooperatively analyze data and contribute to a common knowledge base • Discovery of published (structured) data and its provenance 	<ul style="list-style-type: none"> • 100 Gigabit/sec <ul style="list-style-type: none"> – lambda based point-to-point for single high b/w flows – capacity planning • Network monitoring 	<ul style="list-style-type: none"> • Track world-wide resource usage patterns to maximize utilization • Direct network access to data management systems • Monitoring to enable optimized use of network, compute, and storage resources • Publish / subscribe and global discovery
HEP (5-10 yr)		<ul style="list-style-type: none"> • 1000s of petabytes of data 		<ul style="list-style-type: none"> • 1000 Gigabit/sec 	

Feature	Characteristics that Motivate H-S Nets	Vision for the Future Process of Science	Anticipated Requirements	
Discipline			Networking	Middleware / Grid Services
Chem. Sci. (near term)	<ul style="list-style-type: none"> • High data-rate instruments • Greatly increased simulation resolution- data sets ~10 – 30 TB • Geographically separated resources (compute, viz, storage, instmnts) & people • Numerical fidelity and repeatability • Cataloguing of data from a large number of instruments 	<ul style="list-style-type: none"> • Distributed collaboration • Remote instrument operation / steering • Remote visualization • Sharing of data and metadata using web-based data services 	<ul style="list-style-type: none"> • Robust connectivity • Reliable data transfer • High data-rate, reliable multicast • QoS • International interoperability for namespace, security 	<ul style="list-style-type: none"> • Collaboration infrastructure • Management of metadata • High data integrity • Global event services • Cross discipline repositories • Network caching • Server side data processing • Virtual production to improve tracability of data • Data Grid broker / planner • Cataloguing as a service
Chem. Sci. (5 yr)	<ul style="list-style-type: none"> • 3D Simulation data sets 30 - 100 TB • Coupling of MPP quantum chemistry and molecular dynamics simulations • Validation using large experimental data sets 	<ul style="list-style-type: none"> • Remote steering of simulation time step • Remote data sub-setting, mining, and visualization • Shared data/metadata w annotation evolves to knowledge base 	<ul style="list-style-type: none"> • 10's Gigabit for collaborative viz and mining of large data sets 	<ul style="list-style-type: none"> • Remote I/O • Collaborative • International interoperability for collab. infrastructure, repositories, search, and notification • Archival publication • Knowledge management to facilitate science-level, Grid based problem solving systems (the “Semantic Grid”)
Chem. Sci. (5+ yr)	<ul style="list-style-type: none"> • Accumulation of archived simulation feature data and simulation data sets • Multi-physics and soot simulation data sets ~1 PB 	<ul style="list-style-type: none"> • Internationally collab knowledgebase • Remote collaborative simulation steering, mining, viz. 	<ul style="list-style-type: none"> • 100? Gigabit for distributed computation chemistry and molecular dynamics simulations 	<ul style="list-style-type: none"> • Remote collaborative simulation steering, mining, viz. • Data provenance • Ontologies and constraints for automatic query resolution

Feature	Characteristics that Motivate H-S Nets	Vision for the Future Process of Science	Anticipated Requirements	
Discipline			Networking	Middleware / Grid Services
Magnetic Fusion (near term)	<ul style="list-style-type: none"> • Each experiment only gets a few days per year - high productivity is critical • 100 MBy every 15 minutes to be delivered in two minutes • Highly collaborative environment 	<ul style="list-style-type: none"> • Real time data analysis for experiment steering (the more that you can analyze between shots the more effective you can make the next shot) 		
Fusion (5 yr)	<ul style="list-style-type: none"> • 1000 MBy generated by experiment every 15 minutes (time between shots) to be delivered in two minutes • 1000 MBy generated by simulation to be delivered in two minutes for comparison with experiment • Simulation data scattered across US • Transparent security • Global directory and naming services needed to anchor all of the distributed metadata • Support for “smooth” collaboration in a high stress environments 	<ul style="list-style-type: none"> • Real time data analysis for experiment steering combined with simulation interaction = big productivity increase • Real time visualization and interaction among collaborators across US • Integrated simulation of the several distinct regions of the reactor will produce a much more realistic model 	<ul style="list-style-type: none"> • 500 Mbit/sec for 20 seconds out of 15 minutes, guaranteed • QoS • 5 to 10 remote sites involved for data analysis and visualization 	<ul style="list-style-type: none"> • Parallel network I/O between simulations, data archives, experiments, and visualization • High quality, 7x24 PKI infrastructure • end-to-end QoS • QoS management • Secure / authenticated transport to ease access through firewalls • Reliable data transfer • transient and transparent data replication for real-time reliability • Collaboration support
Fusion (5+ yr)		<ul style="list-style-type: none"> • Real time remote operation of the experiment 	<ul style="list-style-type: none"> • QoS for latency and reliability 	

Evolving Requirements for Network Related Infrastructure

1-3 yrs	 <p>1-40 Gb/s, end-to-end</p> <p> S storage C compute I instrument C&C cache & compute </p>	 <p>guaranteed bandwidth paths</p>
<p>In the near term applications need high bandwidth</p>		<p>2-4 yrs requirement is for high bandwidth and QoS.</p>
3-5 yrs		
<p>3-5 yrs requirement is for high bandwidth and QoS and network resident cache and compute elements.</p>		<p>4-7 yrs requirement is for high bandwidth and QoS and network resident cache and compute elements, and robust bandwidth (multiple paths)</p>

Evolving Requirements for Middleware

Capabilities to support scientists / engineers / domain problem solvers

- Collaboration tools (work group management, document sharing and distributed authoring, sharing application session, human communication)
- Programmable portals - facilities to express, manipulate, preserve the representation of scientific problem solving steps (AVS, MatLab, Excel, SciRun)

- Data discovery (“super SQL” for globally distributed data repositories), management, mining, cataloguing, publish
- Human interfaces (PDAs, Web clients, high-end graphics workstations)
- Tools to build/manage dynamic virtual organizations
- Knowledge management

Capabilities to support building the portals / frameworks / problem solving environments

- Resource discovery, brokering, job management
- Workflow management
- Grid management – fault detection and correction

- Grid monitoring and information distribution – event publish and subscribe
- Security and authorization,

Capabilities to support instantiating science scenarios as computational models

- Utilities for visualization, data management (global naming, location transparency (replication mgmt, caching), metadata management, data duration, discovery mechanisms)

- Support for programming on the Grid (Grid MPI, Globus I/O, Grid debuggers, programming environments, e.g., to support the model coupling frameworks, and to), Grid program execution environment
- User services (documentation, training, evangelism)

Capabilities to support building Grid systems

Cyberinfrastructure for Science

- Such complex and data intensive scenarios require sophisticated, integrated, and high performance infrastructure to support the application frameworks that are needed to successfully manage and carry out the many operations needed to accomplish the science:
 - high-speed networks
 - very high-speed computers
 - ***highly capable middleware (the primary topic of this talk)***

Highly Capable Middleware

- Core Grid services
 - Provide the consistent and uniform foundation for managing dynamic and administratively heterogeneous pools of compute, data, and instrument resources
- Higher level services
 - Provide value-added, complex, and aggregated services to users and application frameworks
- Information management
 - Data Grid services are trying to provide a consistent and versatile view of data of all descriptions
- Knowledge management
 - Services for unifying, classifying and manipulating services, data, and information in the context of a human centric problem solving environment – the Semantic Grid

What are the Core Grid Middleware Services?

- A common security model for Grid services and Grid applications provides uniform and versatile authentication, authorization, and privacy
 - Basis of the cyber-trust that enables collaboration among the many organizations of a large science project
 - Preserves local autonomy of resource owners
 - Correctly used, provides pretty good security
- Standardized access to computing systems and data storage systems
- Tools and services supporting construction and management of collaborations (virtual organizations)

What are the Core Grid Middleware Services

- Services for dynamic construction of execution environments supporting complex distributed applications
 - locating and co-scheduling many resources to support, e.g., transient and complex, science and engineering experiments that require combinations of instruments, compute systems, data archives, and network bandwidth at multiple locations
- Management of dynamic pools of underlying resources
 - Automatic resource registration and de-registration
 - Resource discovery
- Evolving to a Web Services / object oriented model of core services (more later)

Architecture of a Grid

Science Portals

Portal Toolkits (to build application portals)

Applications (Simulations, Data Analysis, etc.)

Grid Services (Workflow management, Visualization, Data Publication/Subscription, Brokering, Job Mg'mt, Data Grid Services, Fault Mg'mt, Grid System Admin., etc.)

Grid Core Functions: Standardized Services and Resources Interfaces

Grid
Information
Service

Uniform
Computing
Access

Uniform Data
Access

Unix and
OGSA hosting

Co-
Scheduling

Global Event
Services,
Auditing,
Monitoring

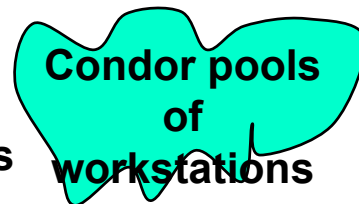
Collaboration
and Remote
Instrument
Services

Systems
management
and access

Communication
Services

Authentication
Authorization

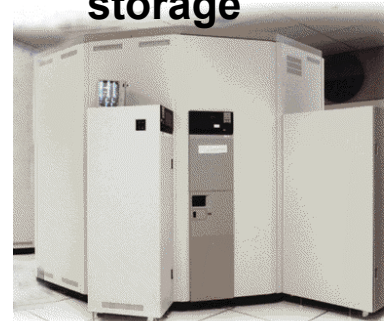
Security
Services



network
caches



tertiary
storage



scientific instruments



High Speed Communication Services



Higher Level Grid Middleware

- Provides standardized services to science portals and applications, e.g. for
 - Distributing and managing massive datasets that must be accessible accessible by world-wide collaborations – e.g. High Energy Physics and Earth Sciences Data Grids
 - Managing complex workflow involving many compute and data intensive steps occurring at different geographic locations
 - Virtual data services for on-demand data generation
 - Autonomous fault management and recovery for both applications and infrastructure
 - And so forth ...

What is Grid Middleware?

- Grids are also several hundred people from the US, European, and SE Asian countries working on best practice and standards at the Global Grid Forum (www.gridforum.org)
- A major industry effort to combine Grid Services and Web Services (IBM, HP, Microsoft)
 - e.g. see “Developing Grid Computing Applications” in IBM developerWorks : Web services : Web services articles - <http://www-106.ibm.com/developerworks/library/ws-grid2/?n-ws-1252>
- Vendor support from dozens of IT companies

The State of Grids

- Persistent infrastructure is being built²
 - Core Grid services are being maintained on the compute and data systems in prototype production Grids
 - Grid security infrastructure supporting “single sign-on” and cross-site trust is operational in all major Grids
 - Resource discovery services are being maintained (Grid Information Service – distributed directory service)
- This is happening, e.g., in
 - NASA’s IPG³
 - DOE Science Grid⁴
 - EU Data Grid⁵
 - UK eScience Grid⁶
 - NSF TeraGrid⁷

all of which are focused on large-scale science and eng.

User Interfaces
Application Frameworks
Applications

Grid Services: Uniform access to distributed resources

Grid Information Service	Uniform Resource Access	Brokering	Global Queuing	Co-Scheduling	Global Event Services	Data Cataloguing	Uniform Data Access	Collaboration and Remote Instrument Services	Network Cache	Communication Services	Authentication Authorization	Security Services	Auditing	Monitoring	Fault Management
--------------------------	-------------------------	-----------	----------------	---------------	-----------------------	------------------	---------------------	--	---------------	------------------------	------------------------------	-------------------	----------	------------	------------------

Grid Managed Resources

DOE Science Grid

Asia-Pacific

PNNL



Europe

ESNet

MDS
CA



NERSC

Supercomputing
& Large-Scale
Storage

ANL



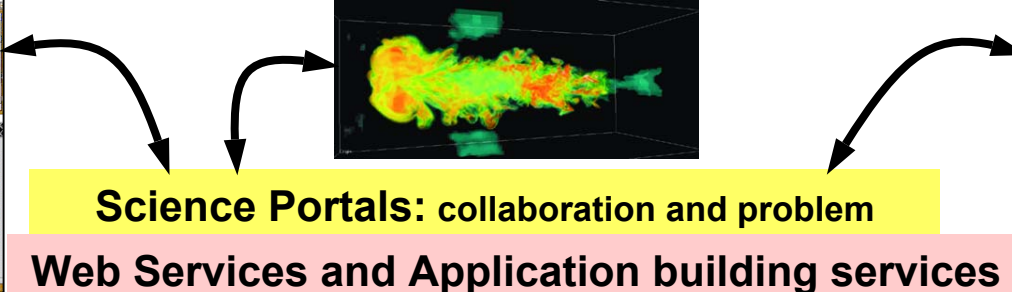
ORNL

LBNL

Initial Science Grid Configuration

Funded by the U.S. Dept. of Energy, Office of Science,
Office of Advanced Scientific Computing Research,

Mathematical, Information, and Computational Sciences Division



Grid Services: secure and uniform access and management for distributed resources

Supercomputing and Large-Scale Storage



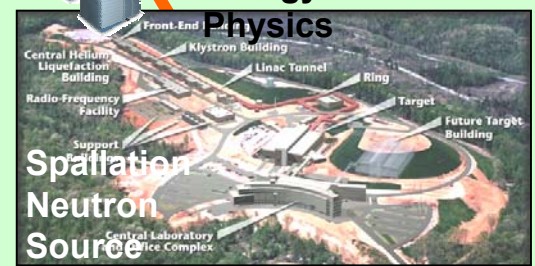
Advanced Engine Design



Macromolecular Crystallography



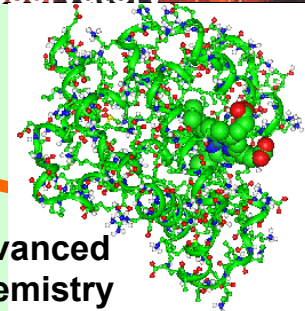
Advanced Photon Source



Spallation Neutron Source



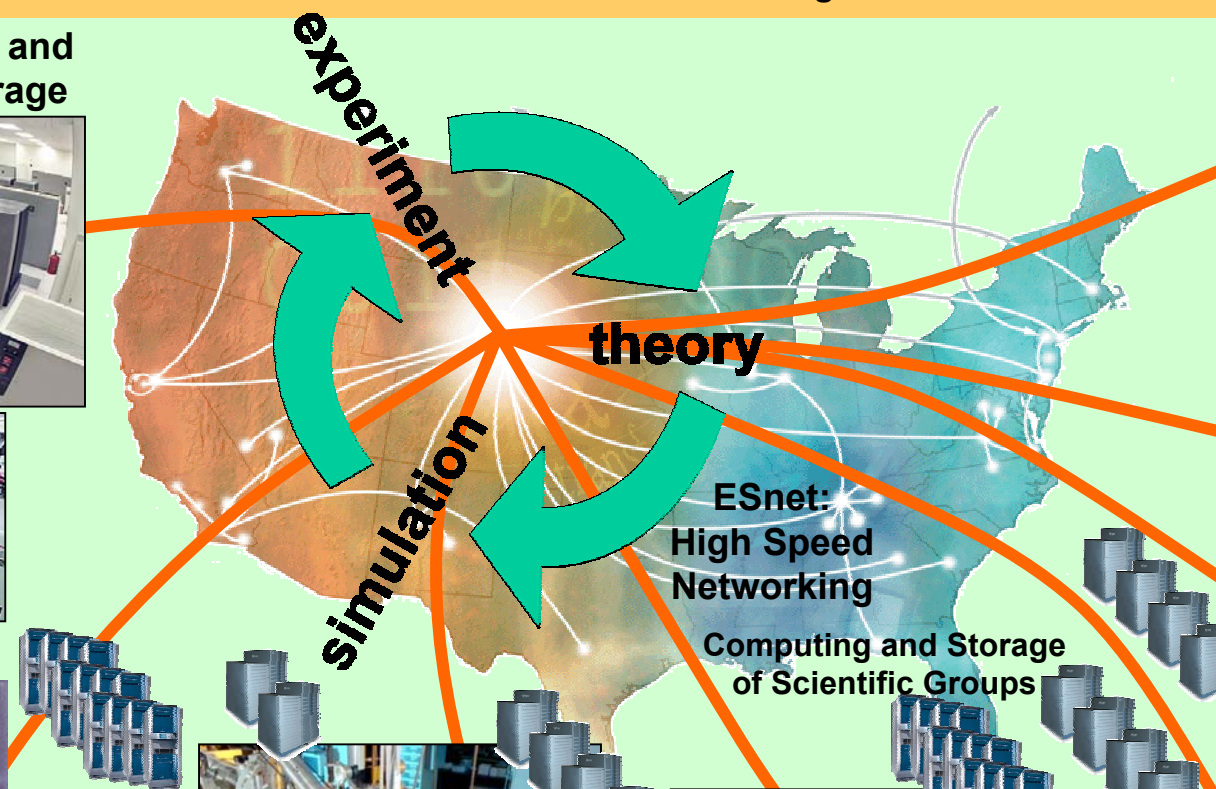
Supernova Observatory



Advanced Chemistry



High Energy Physics



The Evolving Grid - Data Grid Services, Web Services, and the Semantic Grid - is part of the Evolving Cyber Infrastructure:

What follows:

- 1) Current cyber infrastructure
- 2) Evolving Data Grid services as they are currently being experimented with in the High Energy Physics, astronomy, and astrophysics communities.
- 3) Evolution of Cyberinfrastructure and the Grid:
Implications of Combining Web Services and Grids
- 4) The potential of a Semantic Grid / Knowledge Grid:
Combining Semantic Web Services and Grid Services

1) Current Cyberinfrastructure

technology	impact
<p>Internet Protocol (IP): transport independent of the type of the underlying network.</p> <p>TCP: reliable transport</p> <p>Routing: how get a packet to its destination</p> <p>Domain Name System: basic directory service – e.g. lbl.gov ->128.3.7.82)</p> <p>High-speed networks, fairly widely deployed</p>	<p>Internet and physical network technologies have given us a basic global cyberinfrastructure for transport and interprocess communication.</p>
<p>Secure Socket Layer (SSL)/ Transport Layer Security (TLS)</p>	<p>SSL provides the ability to communicate securely between know and authenticated endpoints.</p>
<p>Hyper Text Markup Language (HTML – standardized, low-level document formatting) and the client-server Web.</p>	<p>The Web gave us an information bazaar.</p>
<p>Grids: services for dynamically managing and securely accessing heterogeneous, distributed, and administratively diverse, compute, data, instrument, and human resources</p>	<p>The Grid is providing a global infrastructure for large-scale collaborative science.</p>

2) Evolving Data Grid services in High Energy Physics

- Data Grid services are being developed in the NSF GriPhyN, DOE PPDG, and EU DataGrid projects for the High Energy Physics, astronomy, and astrophysics communities. The Data Grid services include
 - uniform and high-speed access in world wide collaborations
 - data replica catalogues
 - transparency with respect to location
(Data may be replicated at various locations in order to avoid long transmission times (e.g. trans-Atlantic or trans-Pacific). This should be transparent to the users.)
 - virtual data
 - transparency with respect to materialization
(When you request data it might come from storage, or it might be recomputed on the fly, depending on which is quicker.)

Data Replica Services*

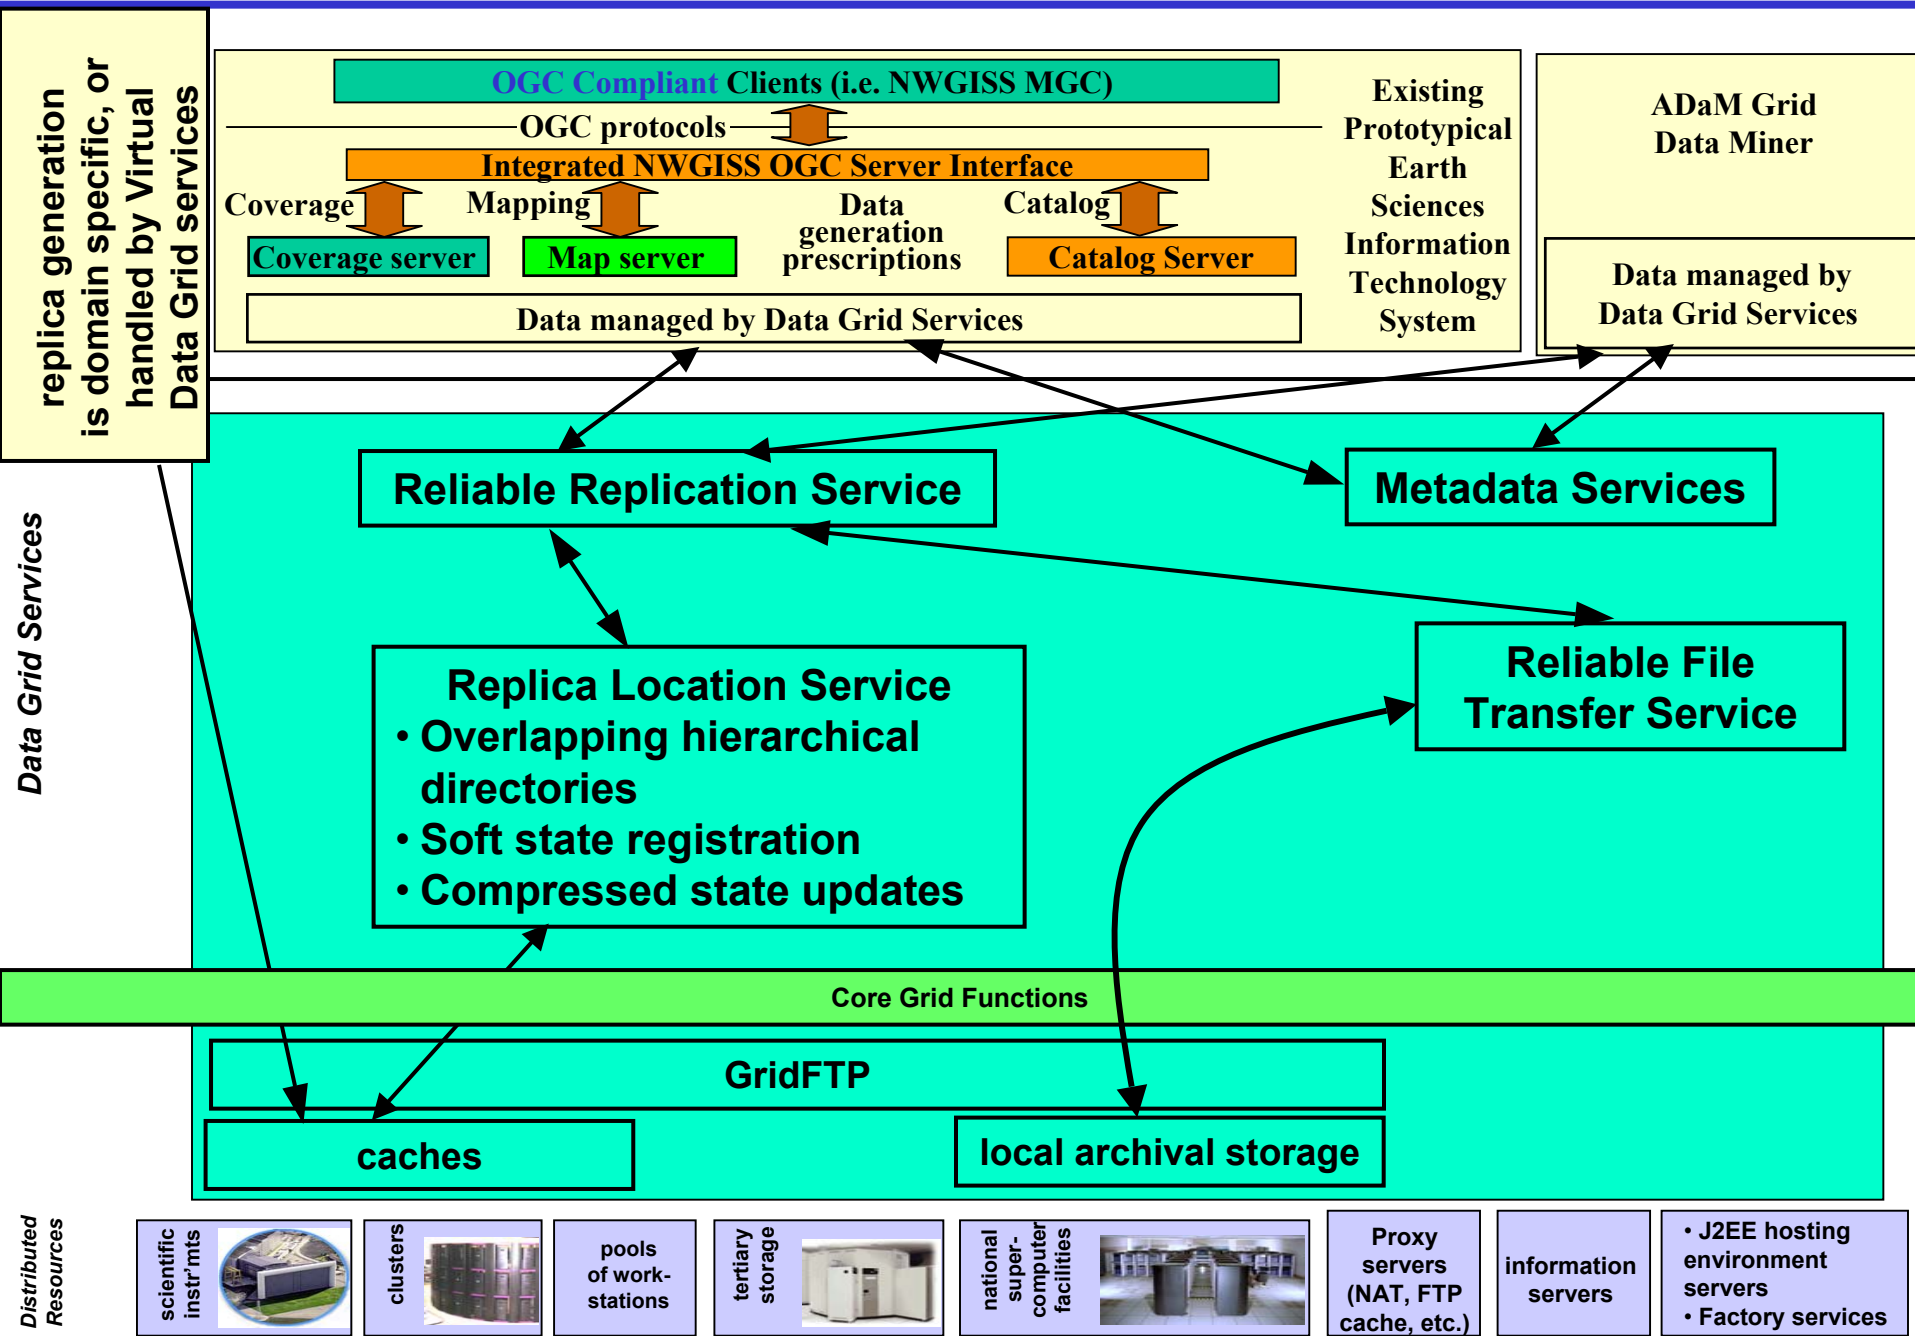
- Naming transparency
 - persistent logical names regardless of physical storage
- Location transparency
 - a logical file name may represent data in several different locations in order to allow for optimizing access by many different users – that is, the service returns the physical file name that is best suited (usually fastest access) for a user
- Data Replica Services include
 - Secure, efficient, high-performance wide area data transfer (e.g. Globus GridFTP)
 - Reliable file transfer service (coordinated, fault tolerant data movement)
 - Reliable physical file replication management
 - Replica Location Service (scalable mechanisms for managing global logical dataset names and data replicas)
 - Metadata service for information that describes logical files

*See: “Giggle: A Framework for Constructing Scalable Replica Location Services” SC02 (<http://www.globus.org/research/papers/giggle.pdf>). Giggle (Gigascale Global Location Engine) is a unified project of NSF GriPhyN, EU DataGrid, and DOE Data Grid Toolkit.

HEP Replica Dataset Management Requirement

- Mostly read-only data and versioning
- Several hundred replica sites, 50 million logical files, 500 million physical files
- 1000 queries and 200 updates per second, average response time ≤ 10 ms, max. query time ≤ 5 s
- Protect privacy and integrity of the metadata and replica catalogues
- A Replica Service need not provide a completely consistent view of all available replicas (however, while this may cause the client to execute less efficiently, it will not execute incorrectly)
- The replica service should be capable of being configured so that there are no single points of failure (this does not address reliability of the primary data storage)

Data Replica Services¹¹



Virtual Data Services*

- “Virtual data:” a class of operations that, for example, “re-materialize” data products
 - that were deleted or generate data products that were defined but never created
 - when data dependencies or transformation programs change
 - for the purpose of creating replicas of data products at remote locations when re-creation is more efficient than data transfer.
- Additionally, virtual data catalogues can track how data products are derived—with sufficient precision so that one can then explain definitively how data products are created, a data provenance aspect that is often not feasible even in carefully curated databases.
This addresses one aspect of overall data provenance.

* See “Chimera: A Virtual Data System for Representing, Querying, and Automating Data Derivation,” I. Foster, J. Voeckler, M. Wilde and Y. Zhao. In *14th Intl. Conf. on Scientific and Statistical Database Management*. 2002. Edinburgh, Scotland.

Background: CMS Data Analysis Scenario¹²

In a typical analysis, a physicist selects events from a large database consisting of small sized event classification information, or TAG data (metadata). Using the TAG data, the physicist gathers the full reconstructed event from various sources including, if necessary the creation of fully reconstructed versions of those events if they do not already exist on some "convenient" storage system. A typical analysis might look something like the following:

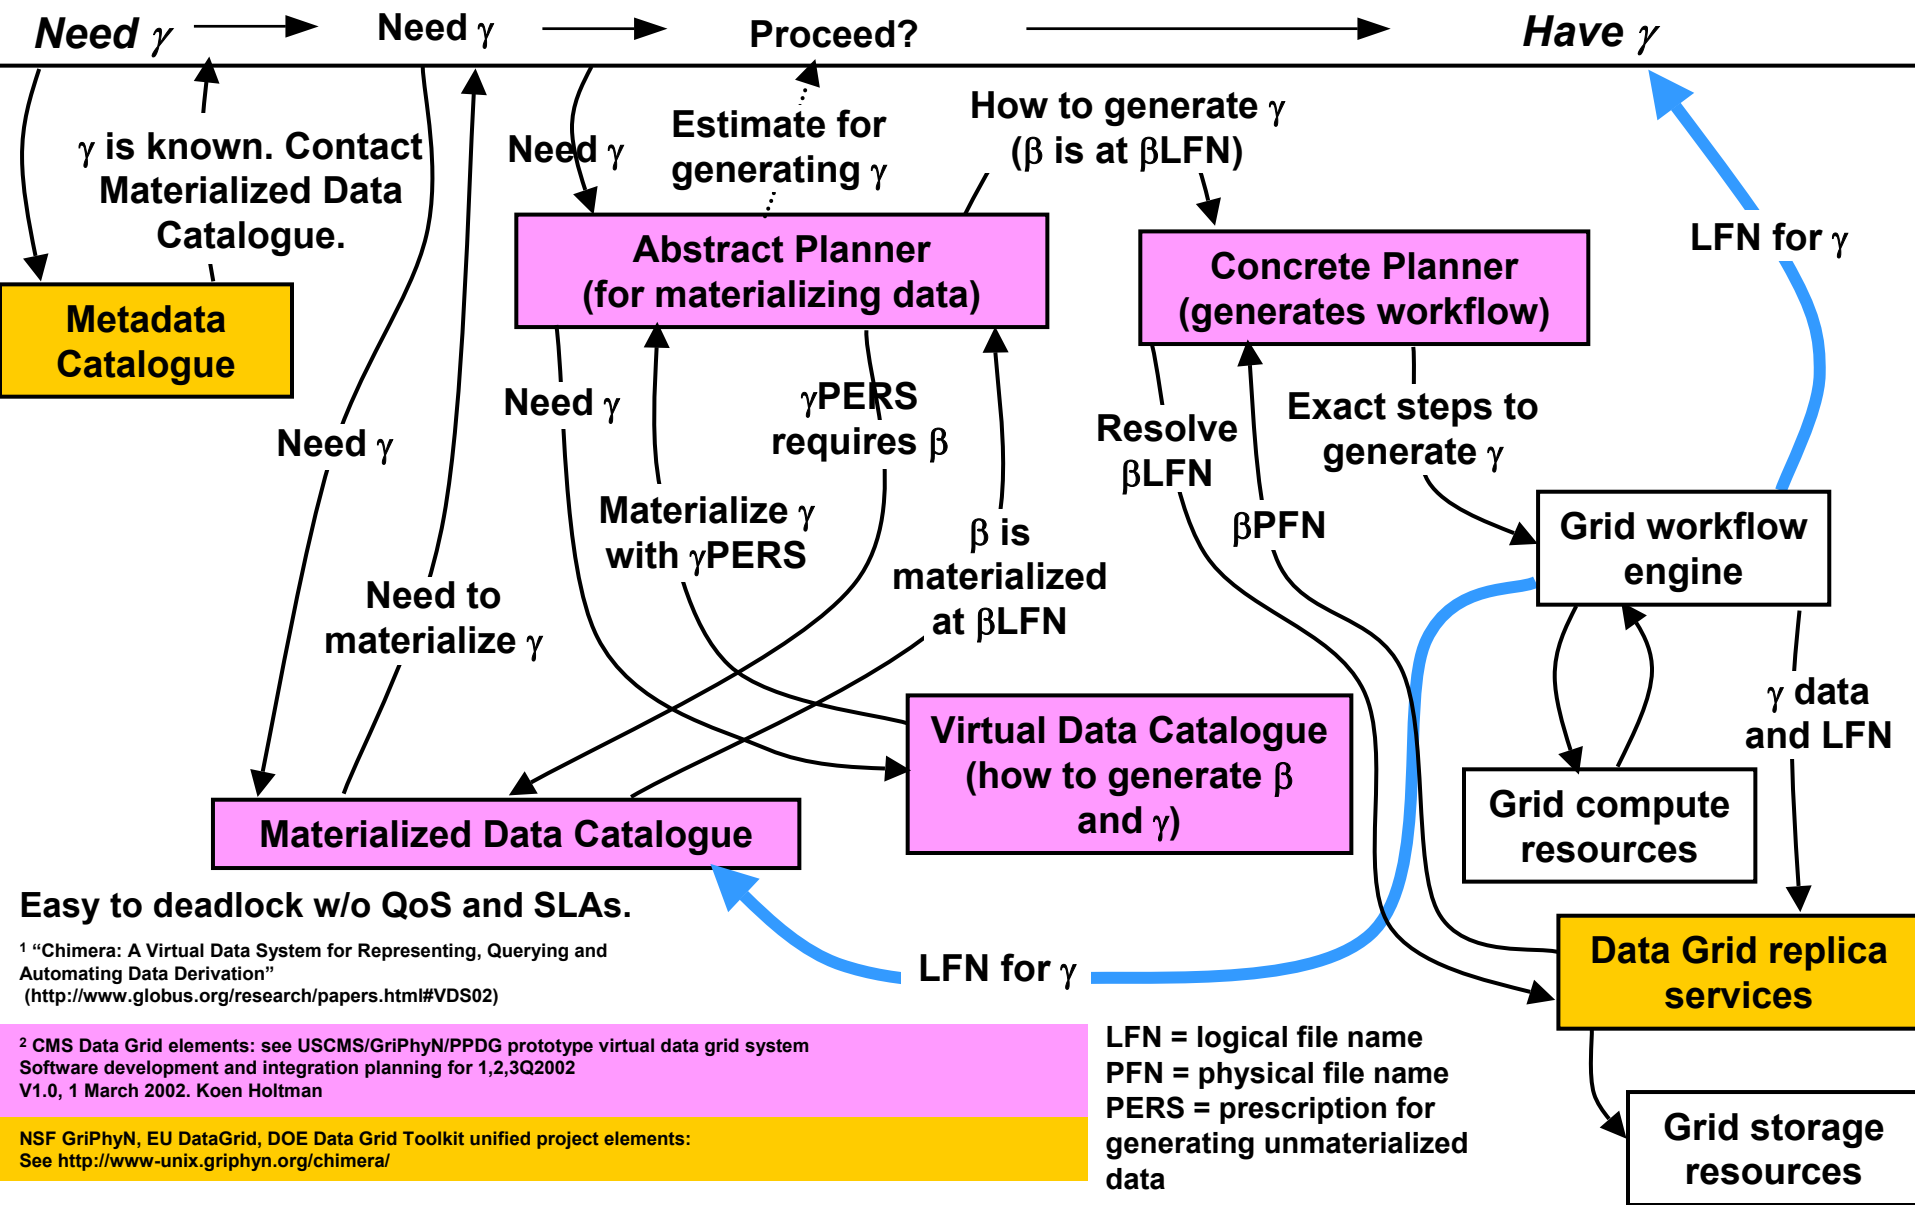
- Search through 109 TAG events (~100 GB), select 106 of those events and fetch the Event Summary Data (ESD) data for the selected events (~1 TB)
- Invoke a user defined reconstruction of the ESD data and make a user defined set of Analysis Object Data (~100 GB) and Tags (~100 MB)
- Analyze the user defined Analysis Object Data (AOD) and TAG datasets interactively, extracting a few hundred candidate signal events (~10 MB).
- Histogram the results and visualize some of the "interesting" events. Virtual data techniques will track data dependencies for the files and/or objects in this process from TAG schemas and TAG databases (or tables) back to the reconstructed event sets and possibly back to the raw data.

Virtual Data Grid Example

- The following example is derived from work done in the High Energy Physics community, but the approach appears to be very general, with potential applicability to
 - defining derived data products and generating them on-the-fly
 - defining simulation parameter spaces for inclusion in data catalogues that are used by downstream applications. The actual parameter studies can be generated on-demand using the Virtual Data Grid mechanisms.

Virtual Data Grid¹ Example (Adapted from High Energy Physics, LHC/CMS Data Grid²)

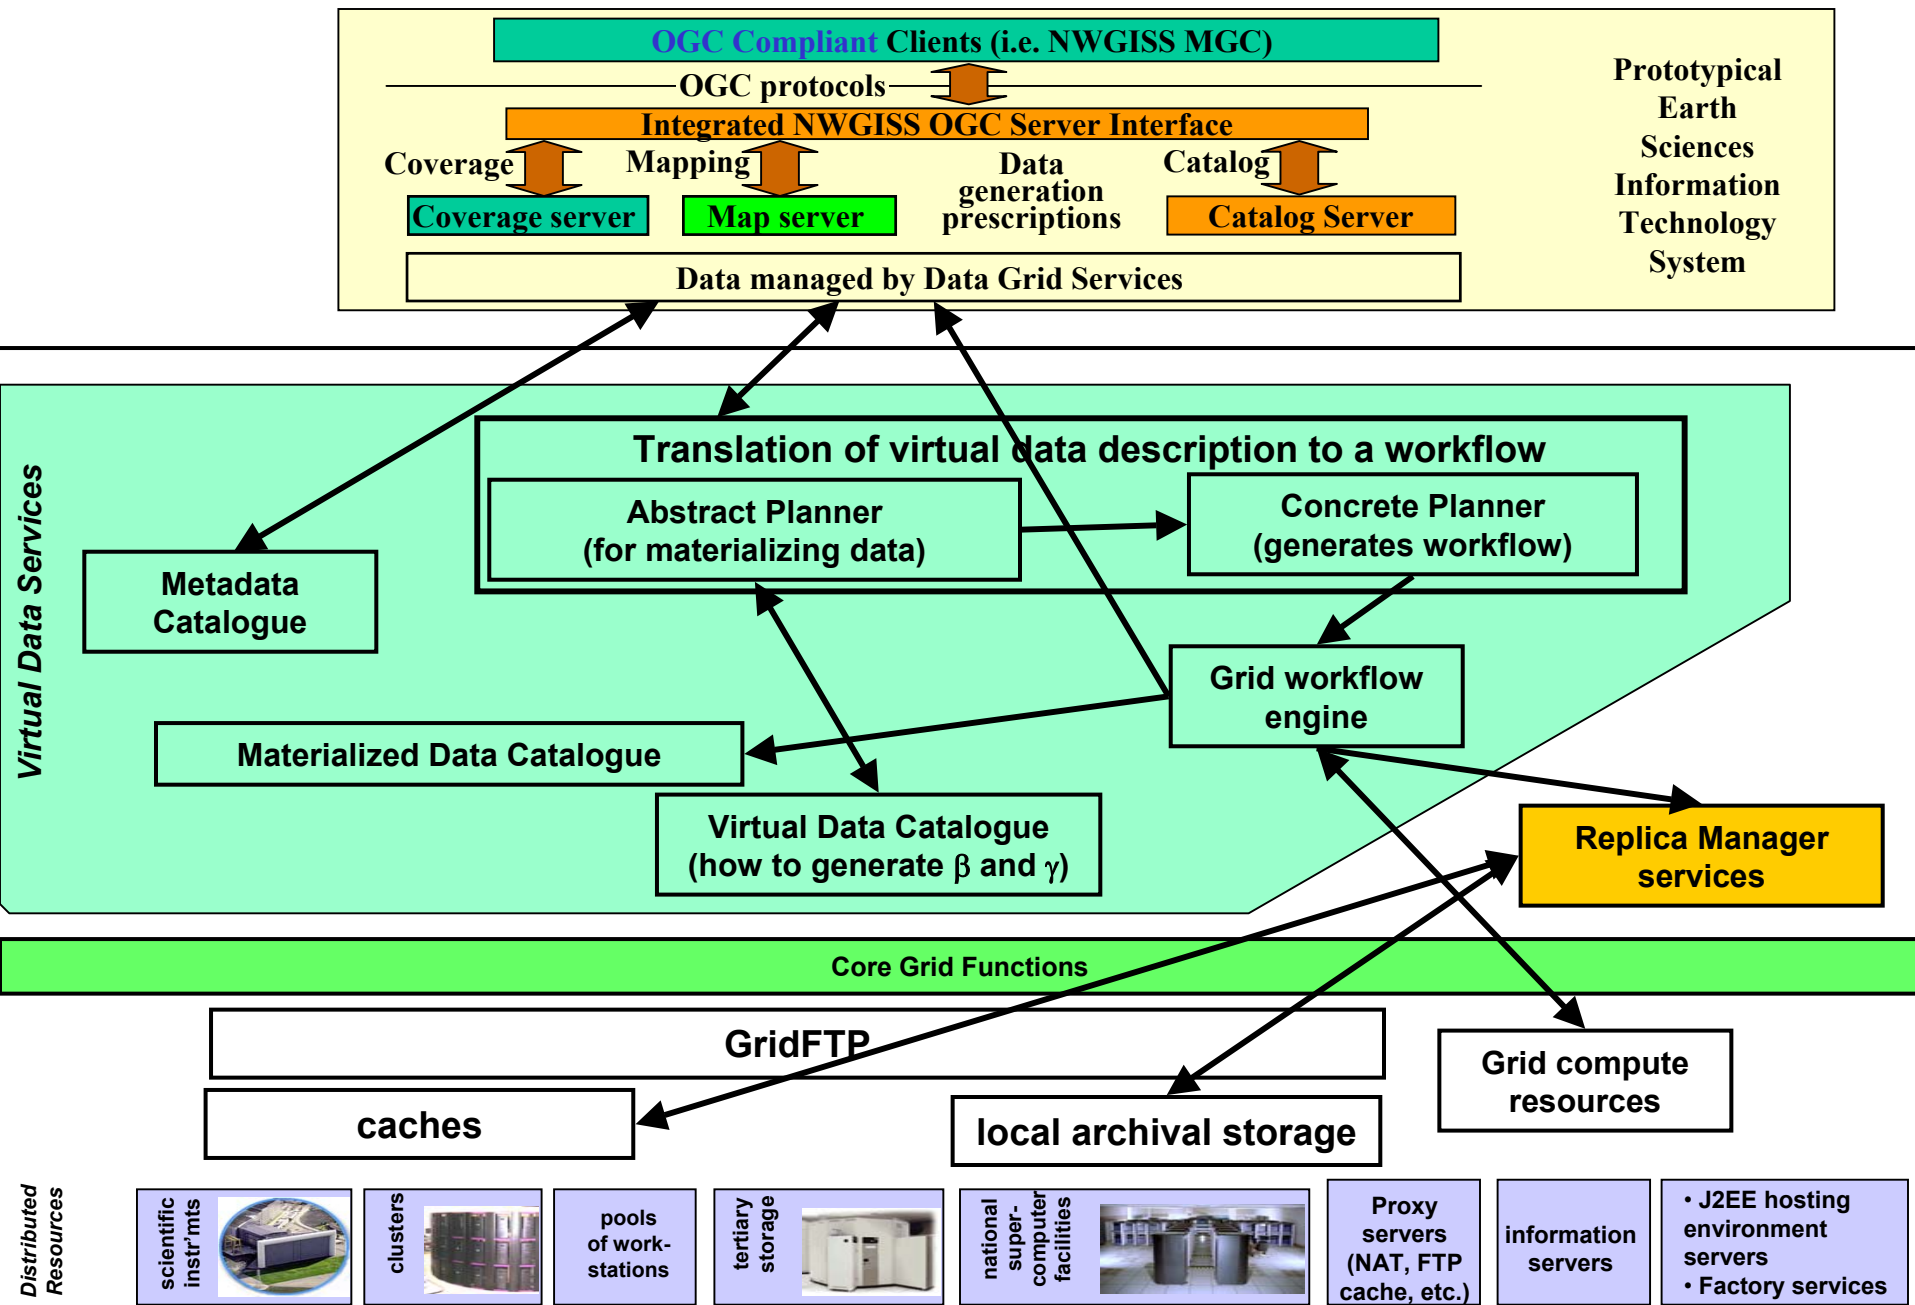
Application: Three data types of interest: γ is derived from β , β is derived from α , which is primary data (interaction and operations proceed left to right)



Virtual Data Grid Components

- This Virtual Data Grid example demonstrates a full range of VDG issues and capabilities
 - Virtual data enabled metadata catalogue that know what data is know, materialized or not
 - Materialized data catalogue that keeps track of what data is actually materialized
 - Virtual data catalogue manages the prescriptions that describe how to re-materialize / regenerate data
 - Abstract planner converts data re-materialization prescriptions into general workflow descriptions
 - Concrete planner converts general workflow steps and logical file names into a specific workflow description that has actual resources and physical files names
 - Workflow engine that executes the re-materialization plan
 - Data Grid replication services manage instances of the data after it is produced

Virtual Data Grid Components



The Virtual Data Grid is Being Prototyped*

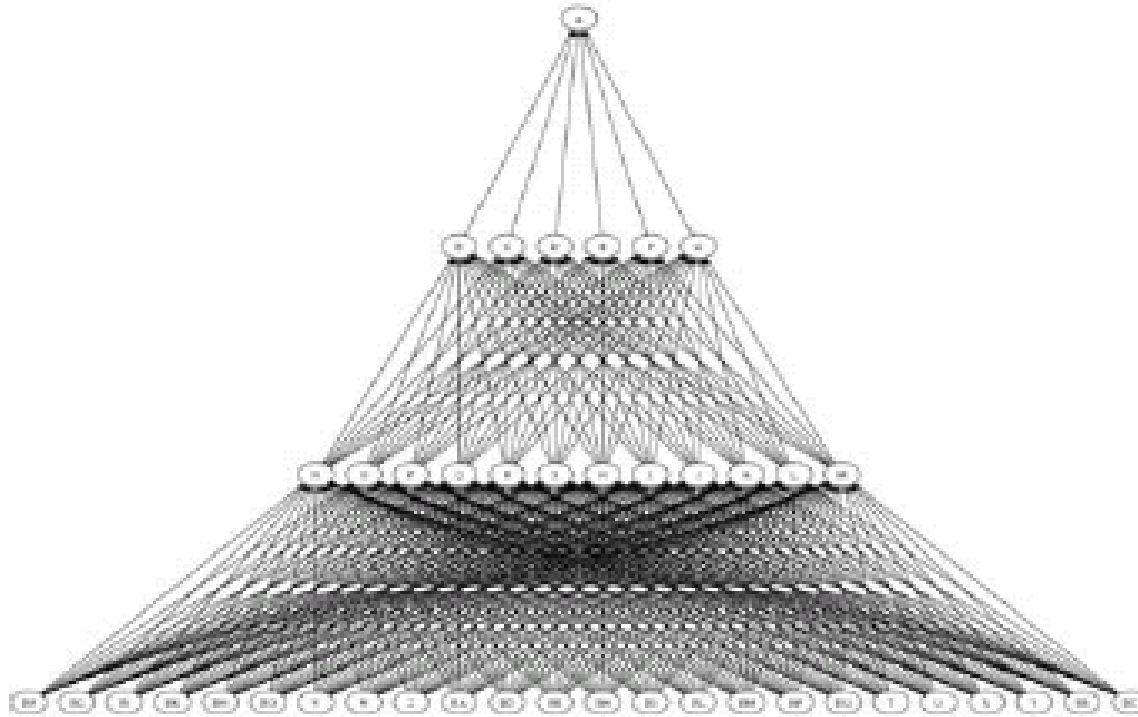


Figure 7: DAG for cluster identification workflow.

- A DAG representation of the workflow generated by a concrete planner represents 48 and 60 searches over 600 datasets
- DAG (each node represents a process on a machine) executed in 2402 seconds on 62 hosts and in 2480 seconds on 48 hosts.

*From "Applying Chimera Virtual Data Concepts to Cluster Finding in the Sloan Sky Survey," J. Annis, Y. Zhao, J. Voekler, M. Wilde, S. Kent and I. Foster. In SC2002. 2002. Baltimore, MD.<http://www.sc2002.org/paperpdfs/pap.pap299.pdf>

3) Evolution of Cyberinfrastructure and the Grid: Implications of Combining Web Services and Grids

technology	impact
eXtended Markup Language (XML) Tagged fields (metadata) and structured “documents” (XML schema)	Web Services – an application of XML – will provide us with a global infrastructure for managing and accessing modular programs (services) and well defined data
Open Grid Services Architecture	Combines Web Services with Computational and Data Grids services to integrate information, analysis tools, and computing and data systems

Web Services and Grids

There is considerable potential benefit to combining Grid Services and Web services.

- Web services provide for
 - Describing services/programs with sufficient information that they can be discovered and used by many people (reusable components)
 - Assembling groups of discovered services into useful problem solving systems
 - Easy integration with scientific databases that use XML based metadata
- Grids provide for accessing and managing compute and data systems, and provide support for Virtual Organizations / collaborations

The Open Grid Services Architecture (Web + Grid Services)

- From Web services
 - Standard interface definition mechanisms
 - Interface and implementation (multiple protocol bindings)
 - local/remote transparency
 - Language interoperability
 - A homogenous architecture basis
- From Grids
 - Service semantics
 - Lifecycle management and transient state
 - Reliability and security models
 - Service discovery
 - Other services: resource management, authorization, etc.
- See <http://www.globus.org/ogsa/>

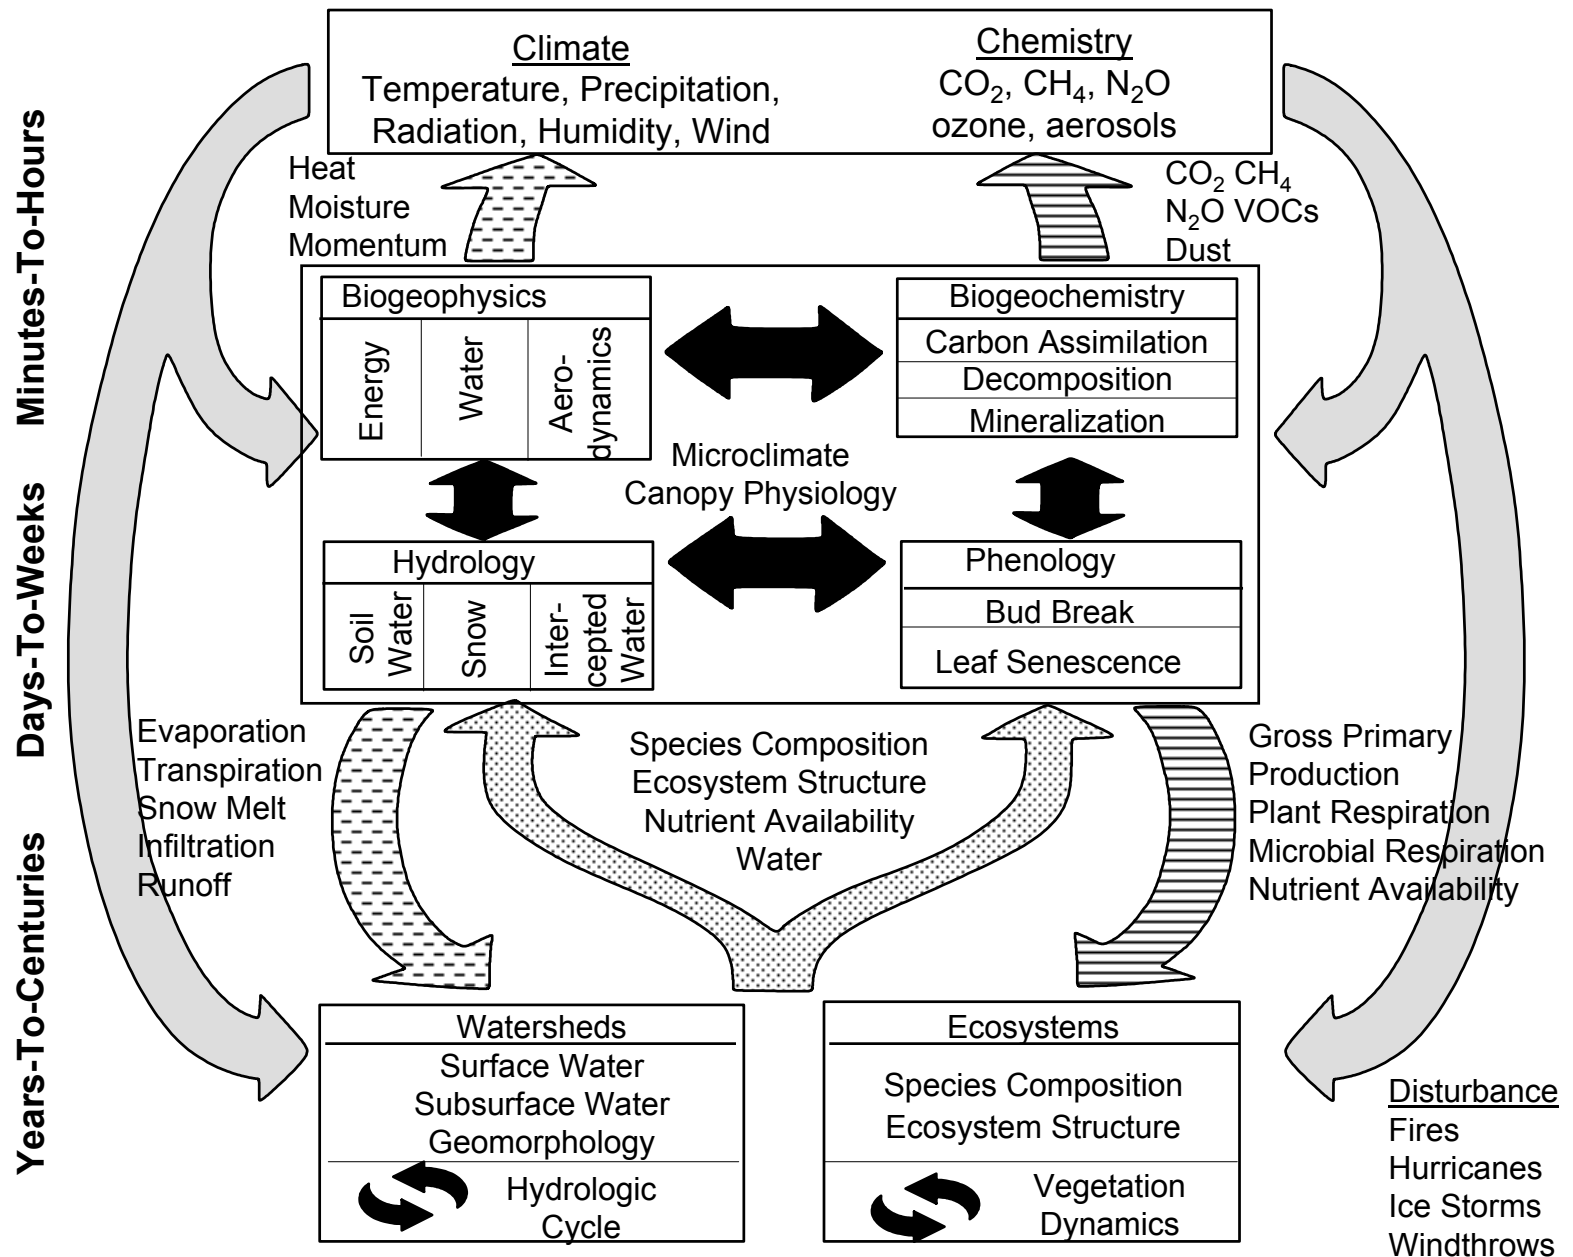
Combining Web Services and Grids

- Combining Grid and Web services provides a dynamic and powerful computing and data environment that is rich in descriptions, services, data, and computing capabilities
- This infrastructure will give us the basic tools to deal with complex, multi-disciplinary, data rich science modeling problems

Combining Web Services and Grids

- Furthermore, the Web Services Description Language, et al, together with Grid services, should be able to provide standardized component descriptions and interface definitions so that compatible sub-models can be “plugged” together
- The complexity of the modeling done in Terrestrial Biogeoscience is a touchstone for this stage of evolution of cyberinfrastructure

Terrestrial Biogeoscience Involves Many Complex Processes and Data



(Courtesy Gordon Bonan, NCAR: *Ecological Climatology: Concepts and Applications*. Cambridge University Press, Cambridge, 2002.)

4) Where to in the Future?

The potential of a Semantic Grid / Knowledge Grid: Combining Semantic Web Services and Grid Services

- Even when we have well integrated Web+Grid services we still do not provide enough structured information to let us ask “what if” questions, and then have the underlying system assemble the required components in a consistent way to answer such a question.

Beyond Web Services and Grids

- A commercial example “what if” question: What does the itinerary look like if I wish to go SFO to Paris, CDG, and then to Bucharest. In Bucharest I want a 3 or 4 star hotel that is within 3 km of the Palace of the Parliament, and the hotel cost may not exceed the U. S. Dept. of State, Foreign Per Diem Rates.
- To answer such a question – relatively easy, but tedious, for a human – the system must “understand” the relationships between maps and locations, between per diem charts and published hotel rates, and it must be able to apply constraints (< 3 km, 3 or 4 star, $\text{cost} < \$ \text{ per diem rates, etc.}$)

Beyond Web Services and Grids

- A science example (necessarily more complex) that is courtesy of Stewart Loken, LBNL is as follows:

HEP experiments collect specific types of data for the particles that result from high energy collisions of the protons, electrons, ions, etc. that are produced by the accelerators. The types of data are a function of the detector and include things like particle charge, mass, energy, 3D trajectory, etc.

However much of science comes from inferring other aspects of the interactions by analyzing what can be observed. Many quantities are used in obtaining the scientific results of the experiment that are derived from what is observed. In doing this more abstract analysis, the physicist typically asks questions like :

Beyond Web Services and Grids

Events of interest are usually characterized by a combination of jets of particles (coming from quark decays) and single particles like electrons and muons. In addition, we look for missing transverse energy (an apparent failure of momentum conservation) that would signal the presence of neutrinos that we cannot detect.

The topologies of individual events follow some statistical distributions so it is really the averages over many events that are of interest. In doing the analysis, we specify what cone angle would characterize a jet, how far one jet needs to be from another (in 3-dimensions), how far from the single particles, how much missing transverse energy, the angles between the missing energy vector and the other particles

What I would like to see is a set of tools to describe these topologies without typing in lots of code. A graphical interface that lets you draw the average event and trace out how statistical variations would affect that. We do simulation of interesting processes and they guide the selection of events, so we would want to learn from that.

In order to automatically transform these sorts of queries into combinations of existing tools and appropriate data queries, ***some sort of knowledge-based framework is needed.***

Knowledge Grids / Semantic Grids

- The emerging Knowledge Grid* / Semantic Grid** services will provide the mechanisms to organize the information and services so that human queries may be correctly structured for the available application services (the model components and data) to build problem solving systems for specific problems
- Work is being adapted from the Artificial Intelligence community to provide
 - Ontology languages to extend terms (metadata) to represent relationships between them
 - Language constructs to express rule based relationships among, and generalizations of the extended terms
- See www.semanticgrid.org and www.isi.cs.cnr.it/kgrid/

* I am indebted to Mario Cannataro, Domenico Talia, and Paolo Trunfio (CNR, Italy), and

** Dave DeRoure (U. Southampton), Carol Gobel (U. Manchester), and Geoff Fox (U. Indiana) for introducing me to these ideas.

Future Cyberinfrastructure*

technology	Impact
<p>Resource Description Framework (RDF)**</p> <p>Expresses relationships among “resources” (URI(L)s) in the form of object-attribute-value (property). Values of can be other resources, thus we can describe arbitrary relationships between multiple resources.</p> <p>RFD uses XML for its syntax.</p>	<p>Can ask questions like “What are a particular property’s permitted values, which types of resources can it describe, and what is its relationship to other properties.”</p>
<p>Resource Description Framework Schema (RDFS)**</p> <p>An extensible, object-oriented type system that effectively represents and defines classes.</p>	<p>Object-oriented structure: Class definitions can be derived from multiple superclasses, and property definitions can specify domain and range constraints.</p> <p>Can now represent tree structured information (e.g. Taxonomies)</p>

* See “The Semantic Web and its Languages,” an edited collection of articles in IEEE Intelligent Systems, Nov. Dec. 2000. D. Fensel, editor.

** The Resource Description Framework,” O. Lassila. *ibid.*

Future Cyberinfrastructure

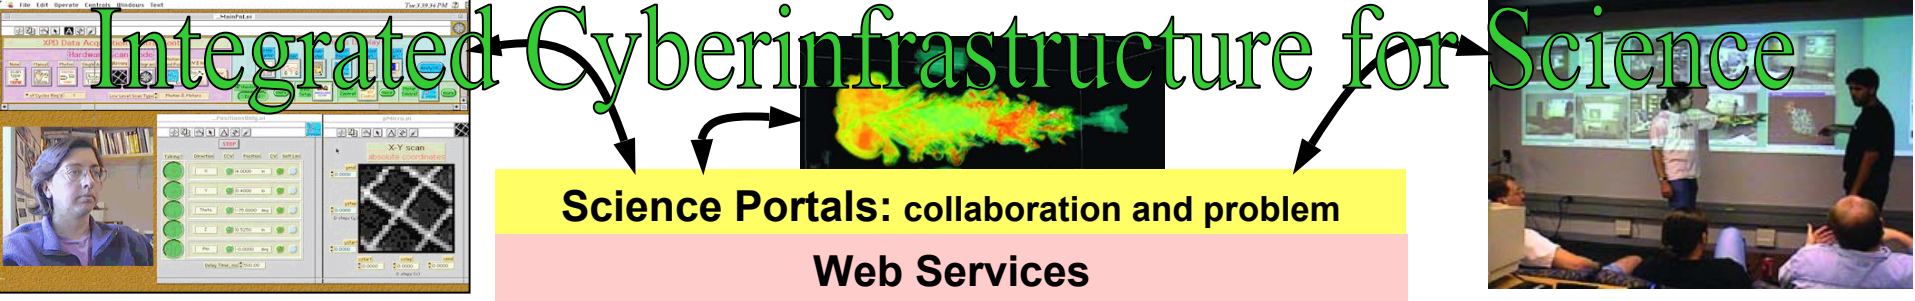
technology	impact
<p>Ontology Inference Layer (OIL)**</p> <p>OIL inherits all of RDFS, and adds expressing class relationships using combinations of intersection (AND), union (OR), and compliment (NOT). Supports concrete data types (integers, strings, etc.)</p>	<p>OIL can state conditions for a class that are both sufficient and necessary. This makes it possible to perform automatic classification: Given a specific object, OIL can automatically decide to which classes the object belongs.</p> <p>This is functionality that should make it possible to ask the sort of constraint and relationship based questions illustrated above.</p>
<p>DAML+OIL+..... ***</p>	<p>Knowledge representation and manipulation that have well defined semantics and representation of constraints and rules for reasoning</p>

** “FAQs on OIL: Ontology Inference Layer,” van Harmelen and Horrocks. *ibid.* and “OIL: An Ontology Infrastructure for the Semantic Web.” *Ibid.*

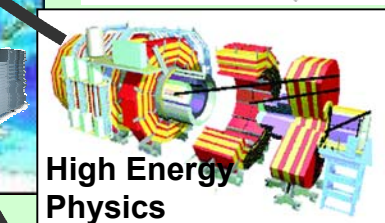
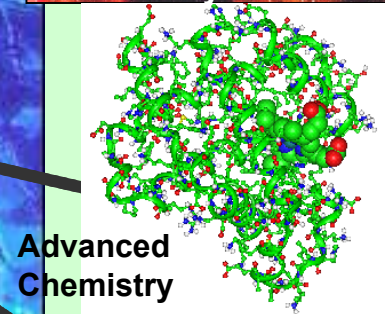
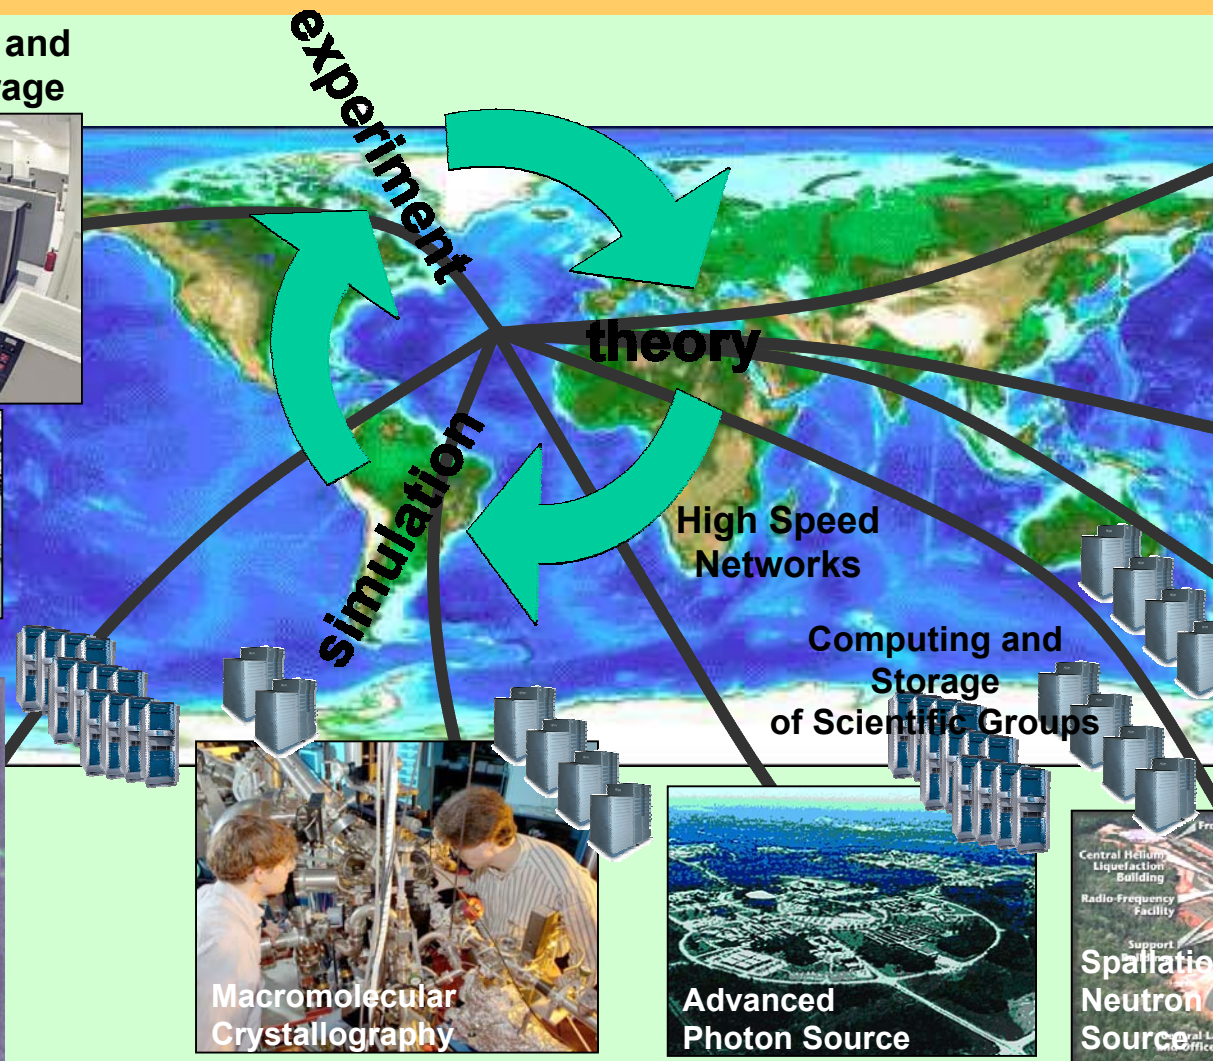
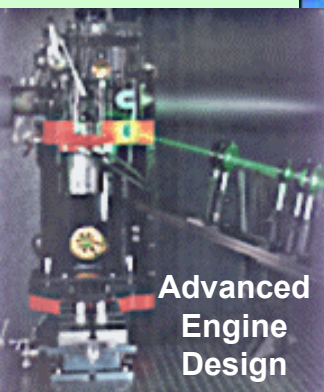
*** “Semantic Web Services,” McIlraith, Son, Zeng. *Ibid.* and “Agents and the Semantic Web,” Hendler. *Ibid.*

Future Cyberinfrastructure

- This Knowledge Grid / Semantic Grid framework should give us the ability to answer “what if” questions by “automatically” structuring data and simulation / analysis components into workflows whose composite actins produce the desired information.
- The Grid Forum has recently established a Semantic Grid Research Group to investigate and report on the path forward for combining Grids and Semantic Web technology.
See <http://www.semanticgrid.org/GGF>
This GGF Research Group is co-chaired by David De Roure <dder@ecs.soton.ac.uk>, Carole Goble <cgoble@cs.man.ac.uk>, and Geoffrey Fox <gcf@grids.ucs.indiana.edu>



Supercomputing and Large-Scale Storage



References

- 1) "The Computing and Data Grid Approach: Infrastructure for Distributed Science Applications," William E. Johnston. <http://www.itg.lbl.gov/~johnston/Grids/homepage.html#CI2002>
- 2) "Implementing Production Grids for Science and Engineering," William E. Johnston, Lawrence Berkeley National Laboratory, Berkeley, Calif. and NASA Ames Research Center, Moffett Field, Calif. (USA); John M. Brooke, Manchester Computing and Department of Computer Science, University of Manchester (UK); Randy Butler, National Center for Supercomputing Applications (NCSA), University of Illinois at Urbana-Champaign (USA); David Foster, CERN LHC Computing Grid Project, Geneva (Switzerland), and; Mirco Mazzucato, INFN-Padova (Italy)
<http://www.itg.lbl.gov/~johnston/Grids/homepage.html#Grid2ed>
- 3) NASA's Information Power Grid - www.ipg.nasa.gov
- 4) DOE Science Grid - www.doesciencegrid.org
- 5) European Union DataGrid Project - www.eu-datagrid.org/
- 6) UK eScience Program - www.research-councils.ac.uk/escience/
- 7) NSF TeraGrid - www.teragrid.org/
- 8) GriPhyN (Grid Physics Network) - <http://www.griphyn.org>
- 9) Particle Physics Data Grid, PPDG. <http://www.ppdg.net/>
- 10) European Union DataGrid Project - www.eu-datagrid.org/
- 11) See: "Giggle: Framework for Constructing Scalable Replica Location Services." Chervenak, et al.
<http://www.globus.org/research/papers/giggle.pdf>
- 12) See "CMS Grid Implementation Plan - 2002," CMS NOTE 2002/015